AMENDMENTS TO THE CLAIMS

Please amend the claims as follows.

1. (Currently Amended) A system comprising:

a plurality of nodes housed within a single computer having a plurality of processors ~~located in a single multiprocessor system~~; and

a mesh interconnect, located within the single computer, connecting the plurality of nodes,

wherein a first node selected from the plurality of nodes comprises a first router for interfacing with the plurality of nodes using the mesh interconnect and a first replicated service executing on a first operating system of the first node,

wherein a second node selected from the plurality of nodes comprises a second router for interfacing with the plurality of nodes using the mesh interconnect and a second replicated service executing on a second operating system of the second node, and

wherein the first node is configured to:

generate, in response to the first service being unavailable, a request to replace the first replicated service ~~when the first replicated service is unavailable~~,

send the request to replace the first replicated service to the plurality of nodes using the mesh interconnect,

receive a response to the request from the second node indicating [[that]] the second node comprises a replacement for the first replicated service,

receive, after receiving the response from the second node, a request for the first replicated service from a third node of the plurality of nodes, and

route, based on the response and using a master-less routing policy implemented by the first router, [[a]] the request for the first replicated service from [[a]] the third node ~~of the plurality of nodes~~ to the second node,

wherein the plurality of nodes comprises a first subset of nodes and a second subset of nodes, wherein the first node is in the first subset and the second node is in the second subset, and wherein the first node is configured to send the request to the second subset of nodes only when the first subset of nodes cannot replace the first replicated service.

2. (Cancelled)

3. (Previously Presented) The system of claim 1, wherein the second node comprises a cache indicating that the second replicated service is available, and wherein the second node is configured to generate the response based on the cache.

4. (Original) The system of claim 1, wherein the first router comprises a lightweight communications protocol.

5. (Previously Presented) The system of claim 1, wherein the first router comprises a heavy-weight communications protocol.

6. (Original) The system of claim 1, wherein the mesh interconnect provides at least two connection paths from the first node to the second node.

7. (Previously Presented) The system of claim 1, wherein the first replicated service is a different application than the second replicated service.

8. (Cancelled)

9. (Cancelled)

10. (Previously Presented) The system of claim 1, wherein the first node is configured to send the first request using at least one selected from a group consisting of a broadcast message and a multicast message.

11. — 25. (Cancelled)

26. (Previously Presented) The system of claim 3, wherein the cache comprises a table having entries for each replicated service provided by the second node.

27. (Previously Presented) The system of claim 1, wherein the first replicated service is unavailable when the first replicated service is busy.

28. (Previously Presented) The system of claim 1, wherein the first replicated service is unavailable when the first replicated service has failed.

29. (Cancelled)

30. (Previously Presented) The system of claim 28, wherein the first replicated service has failed due to a security hole being exploited by a hacker, and wherein the second replicated service does not include the security hole.

31. (Previously Presented) The system of claim 1, wherein the first operating system is different than the second operating system.

32. (Cancelled)

33. (Currently Amended) A method for managing replicated services, comprising:

generating, by a first node selected from a plurality of nodes, a request to replace a first replicated service of the first node when the first replicated service is unavailable, wherein the plurality of nodes is housed within a single computer having a plurality of processors ~~located in a single multiprocessor system~~ and connected using a mesh interconnect;

sending, by the first node, the request to the plurality of nodes using the mesh interconnect;

receiving, at the first node, a response from a second node selected from the plurality of nodes indicating [[that]] the second node comprises a replacement for the first replicated service;

receiving, at the first node and after receiving the response from the second node, a request from a third node of the plurality of nodes for the first replicated service; and

routing, at the first node using a master-less routing policy implemented by a router of the first node, [[a]] the request from [[a]] the third node ~~of the plurality of nodes for the first replicated service~~ to the second node based on the response,

wherein the plurality of nodes comprises a first subset of nodes and a second subset of nodes, wherein the first node is in the first subset and the second node is in the second subset, and where sending the request to the plurality of nodes comprises:

sending the request to the first subset of nodes; and

sending the request to the second subset of nodes when the first subset of nodes cannot replace the first replicated service.

34. (Cancelled)

35. (Currently Amended) A system comprising:

a first node comprising a first router, a first application executing on a first operating system for performing a service, and a cache table having an entry indicating an availability of the service on the first node;

a second node comprising a second router and a second application executing on a second operating system for performing the service, wherein the second node is configured to send a request to replace the service to the first node after failure of the second application; and

a mesh interconnect connecting a plurality of nodes including the first node and the second node,

wherein the first node is configured to examine the entry in the cache based on the request to replace the service, and send a response to the second node using the mesh interconnect,

wherein the second node is configured to:

receive the response from the first node,

receive a request for the service from a third node after receiving the response from the first node, and

route, based on the response and using a master-less routing policy implemented by the second router, [[a]] the request for the service from [[a]] the third node to the first node,

wherein the first node, the second node, the third node, and the mesh interconnect are housed within a single computer having a plurality of processors, and

wherein the first application is different than the second application.

36. (Previously Presented) The system of claim 35, wherein the plurality of nodes comprises a first subset of nodes and a second subset of nodes, wherein the first node is in the first subset and the second node is in the second subset, wherein the second node is configured to send the request to the first subset of nodes only when the second subset of nodes cannot provide

the service, and wherein each of the first subset of nodes and the second subset of nodes includes at least three nodes.

37. (Cancelled)

38. (Previously Presented) The system of claim 35, wherein the second application has failed due to a security hole being exploited by a hacker, and wherein the first application does not include the security hole.